Intel Developer Update is Intel's monthly online news magazine for developers. As the official publication of developer.intel.com, it brings hardware, software, and Web developers the latest information on Intel initiatives, technologies, platforms, and products.

Cover Story

Each month, we run a cover story on the most significant industry announcement, trend, or development for the month.

Featured Articles

Delivering in-depth reports on key platforms, products and technologies, our featured articles provide a monthly source of information on issues affecting developers. Be sure to check in every month for the latest developments driving the evolution of the industry.

Contact the Editor

To make *Intel Developer Update* a better information resource, we invite you to share your thoughts on what we've published or what you'd like to see covered. Comments are always welcome.

Archives

Our archives contain two groups of previously published articles. One group contains all the articles that appeared in *Platform Solutions News*, the earlier version of *Intel Developer Update*. The articles date from September 1997 through August 1999. The other group is set up to contain *Intel Developer Update* articles dating from the inaugural September/October 1999 issue.

Bookmarking

We advise against bookmarking article pages. They're accessible online only during the month the issue is live. Thereafter, they're removed to our archives. Instead, we suggest that you bookmark the PDF (Adobe* Portable Document Format) file versions of the articles. You'll find buttons for the PDF files labeled "print article" in the right navigation section of each article. A PDF for the entire issue is labeled "print magazine" and is located near top right side of the IDU home page.

Table of Contents

(Click on page number to jump to articles)

Note: Intel does not control the content on other company's Web sites or endorse other companies supplying products or services. Any links that take you off of Intel's Web site are provided for your convenience.

Cover Story

Designing for Voice/Data Convergence in the Local Loop

Nabil G. Damouny
Director, Strategic Marketing
Access and Switching Processor Division
Intel Corporation

Overview

The local loop covers the "last mile" or the portion of the Internet between customer premises equipment (CPE) and the local access link. The Telecommunications Act of 1996 requires unbundling of service provisioning in the local loop, and this has created significant new opportunities for carriers, service providers, and the equipment manufacturers who serve them.

In today's hotly competitive broadband market segment, service providers have five major requirements. These include finding new sources of revenue, developing the ability to offer bundled services, finding new ways to enhance customer retention, providing a smooth upgrade path to new service offerings, and finally, discovering innovative ways to reduce provisioning costs.

To meet these requirements, CPE should be able take advantage of existing network infrastructures while meeting ever-higher broadband performance requirements. New equipment should also preserve investments in the existing network infrastructure, while guaranteeing interoperability with transport protocols and delivering wire-speed performance.

The Intel® Internet Exchange Architecture (Intel® IXA) provides building blocks designed to speed time-to-market development of Integrated Access Devices (IAD) for voice and data over digital subscriber line (DSL) technology. The Intel® building block family under Intel IXA consists of silicon, software, and reference designs that enable CPE products to meet specific wide-area network service requirements. This includes the ability to support multiple derived voice channels over the DSL loop, while preserving existing infrastructure, ensuring interoperability, allowing for easy service provisioning, and supporting wire-speed performance.

DSL market drivers

Two forces are driving the deployment and evolution of DSL services. The first driver is the continuing customer demand for high-speed Internet access, provided by cable and DSL services. A second emerging market driver is the convergence of voice and data in the local loop. This convergence has led to the new category of IAD products. Unlike a traditional router or residential gateway appliance which are designed for data only, an IAD is capable of handling not only data but also voice.

An example of the value of an IAD, consider the fact that the simple task of adding a second residential voice line can involve multiple on-site visits by a telephone company service technician. Because each of these "truck rolls" can cost the provider several hundred dollars, the payback period for the added service can be extremely lengthy.

Contrast this with the promise of voice and data capability using an IAD over a DSL line. The capability would allow carriers to provide one, two, four, eight, or even 16 additional derived voice channels over the same DSL copper loop, simply with a flip of a switch. In this way, deployment of voice-and-data DSL IADs and additional equipment at the service provider creates an instant connection to new sources of revenue for a carrier or service provider, using existing copper telephone lines.

Service-specific network processors

The Intel IXA network processor family has recently expanded to include a new generation of service-specific network processors. The Intel® IXP220 and IXP225 Network Processors are the first service-specific network processors optimized to help OEMs quickly meet the CPE requirements of carriers and service providers in the DSL market segment (Figure 1).

Figure 1. The Intel® IXP225 network processor is voice-and-data capable. The Intel® IXP220 network processor is designed for data-only CPE.

These network processors are the heart of a comprehensive IAD design solution that includes hardware and software reference platforms for specific DSL WAN services. The IXP225 network processor is voice-and-data capable, and the IXP220 network processor is designed for data-only CPE, such as bridges, routers, and residential gateways.

Intel's service-specific network processors integrates hardware with software, including a real-time operating system (RTOS), service-specific networking protocol stacks, and an open application programming interface (API) to aid the CPE designer in the easy development of features that differentiate the final product. IXP220 software is a subset of the software stack for the IXP225, which provides developers with a seamless migration path from data-only to converged voice and data solutions.

Intel® IXP225 reference platforms provide proven development frameworks to enable the fast time-to-market development of voice and data applications based on two variants of DSL:

- Asymmetrical DSL (ADSL), which is widely deployed in homes and small offices
- Symmetric high bit-rate DSL (G.SHDSL) for small-to-medium enterprise (SME) environments

Each reference platform provides an IAD solution that supports the convergence of voice and data from telephone and LAN sources to a single CPE WAN port.

Time-to-market

The Intel IXP225 DSL reference platform is available to jump-start the product development cycle. It provides a comprehensive development environment and evaluation platform for ADSL and G.SHDSL IAD products. Integrated voice and data functions include:

- Direct connection for voice and data
- Toll-quality voice
- AAL-2 (ATM Adaptation Layer Type 2) voice processing of variable bit rate (VBR) isochronous traffic, such as compressed voice and data
- Support for multiple voice channels
- Distributed processing architecture, including an ARM7* RISC processor core and two network processor engines (NPE)
- An open software architecture that includes a set of easy to use APIs

The platform routes telephone voice traffic received from analog voice ports or data traffic received through a 10/100 Ethernet PHY to either an ADSL PHY or a G.SHDSL PHY.

The hardware configuration of the Intel IXP225 DSL reference platform includes the Intel IXP225 network processor, the Intel® LXT971 Ethernet PHY, Intel® 28F64OJA flash memory, two dual SLIC/CODECs for four analog voice lines, and either a G.SHDSL PHY or an ADSL PHY. All hardware components provided on the reference platform have been tested and validated to work together for optimized performance and service deployment interoperability.

In addition to their time-to-market benefits, the reference platforms enable measurement of throughput rates, data flow control, and protocol processing performance, while providing a framework for easier platform testing and debugging.

Distributed processing architecture

The distributed processing architecture of the Intel IXP220 and IXP225 network processors features a high-speed core and two dedicated network processor engines. Each network processor chip integrates an ARM7* TDMI RISC processor, a LAN NPE, and a voice/WAN NPE. The function of the NPEs is to offload the processing intensive data manipulation functions from the ARM7* processor. This reduces the burden on the ARM7* CPU and provides the OEM with the headroom to differentiate the IAD by adding intellectual property or existing higher-layer applications that differentiate the IAD

The LAN NPE is associated with the Ethernet interface and provides the processing power needed to support Ethernet filtering at 100 Mbps. Additionally, when it is used for DES3 processing, this NPE provides support for encrypted Ethernet data at 10 Mbps.

The voice/WAN NPE provides support for the ATM functions including the segmentation and re-assembly (SAR) function and the processing required to support ATM adaptation layers 2 and 5 (AAL-2 and AAL-5). This NPE also supports two high-speed serial port network interfaces. The first interface for pulse code modulation (PCM, G.711) or compressed voice (G.728, G.729) inputs and one for connecting to serial WAN PHYs. In addition, this NPE provides the required support for the Utopia WAN port (for ADSL) and the High-speed serial port (for G.SHDSL).

Summary

The DSL CPE market segment continues to grow. In addition to customer demand for high-speed Internet access, the key market driver is the convergence of voice and data in the local loop. Carriers and service providers require a new generation of differentiated IAD products in order to generate new sources of revenue through the flexible and cost-effective provisioning of broadband services to millions of residential, small office/home office (SOHO), and SME customers.

While the opportunities are great and are happening today, this market segment is extremely competitive, and design windows can be short. Intel® IXA silicon, software, and reference platform building blocks are available now to help OEMs address the convergence of voice and data over DSL in the local loop. Intel IXP220 and IXP225 service-specific network processors are the foundations for quick time-to-market development solutions, while providing developers with a fast, cost-effective migration path from data-only to voice-and-data CPE.

More info

Visit the Web site for additional information on the Intel IXP220 and IXP225 Service-Specific Network processors, including product briefs, data sheets, manuals, and a programmer’s guide.

Information on other Intel Networking/Communications Building Blocks is also available online.

For details on Intel® Internet Exchange Architecture, visit the Intel® IXA home page on the Web.

Author Bio

Nabil G. Damouny, Director of Strategic Marketing at Intel's Access and Switching Processor Division (ASPD), has over 20 years experience in communications and networking. Prior to joining Intel, Nabil was a founder and vice president at Basis Communications, a leading provider of integrated networking software and silicon solutions for voice and data wide area networking. Nabil has participated in many standards activities including ISDN, 10Base-T, FDDI, and the ATM-Forum and currently holds three patents in computer architecture and remote networking. He holds B.S.E.E. and M.S.E.C.E degrees.

Column

From the Editor
Donna Loveland
Managing Editor
Intel Developer Update Magazine
Intel Corporation

Column

The Spring IDF Conference in the U.S. is a great place for developers to meet and exchange information about ongoing advances. As thousands of developers took part in labs and sessions and demos on-site in San Jose last month, thousands more were at work implementing the technology and preparing the next wave of innovations.

In this month's issue you'll find seven wide-ranging articles in the areas of applied and desktop computing, initiatives and technologies, networking and communications, and servers.

Designing for Voice/Data Convergence in the Local Loop–cover story–The Intel® IXP225 Reference Platform is a time-to-market design solution that enables OEMs to capture new opportunities driven by the convergence of voice and data services in the local loop.

Keeping Transients and Static in Check–Limiting transients and static in a desktop PC motherboard helps prevent electrical anomalies and protect the sensitive transistors in the CPU. A number of elements help reduce transients and static while complying with the voltage regulator module specification.

Achieving High Availability on Intel® Architecture–As dependence on complex hardware and software systems increases, so does the need for systems with virtually no downtime. The HA Forum is developing open standard building blocks for highly available systems.

Unisys Takes Open Systems into the Data Center–With up to 32 processors, Cellular MultiProcessing (CMP) architecture provides near-linear scalability. It establishes an Intel® processor-based high-performance, high-availability system capable of supporting large-scale e-Commerce and mission-critical applications.

Dual Load Switching–In today's market for ACPI-compliant desktop systems, proper voltage regulator switching implementation can help differentiate a robust motherboard. Intel's dual load switch resolves power distribution problems in motherboard design.

Using the CNR with Multiple Chipset Suppliers–The Communication and Networking Riser (CNR) will give OEMs and system integrators a cost-effective and flexible way to implement new and existing technologies on a single riser card.

New Intel® Desktop Boards Provide More Versatility, Less Cost–The new, full ATX-sized Intel® desktop board D815EEA2 and the micro-ATX Intel® desktop board D815EFV are the second generation of boards based on the proven Intel® 815E chipset and the new Intel® 815EP chipset.

Next month, be sure to visit *Intel Developer Update* for continuing in-depth information for developers.

Enjoy.

Author Bio

Donna Loveland is the editor of *Intel Developer Update* magazine. She joined Intel's Technology and Initiatives Marketing group in 1999 as the editor of Platform Solutions News. Donna began her career with Intel in 1982 as a technical editor in an advanced microprocessor development group. Since then, she's held communications positions in leading-edge technology areas ranging from stereoscopic display to digital broadcast to scalable online content. Donna has a B.A. degree in English from the University of Rochester and an M.A. in expository writing from the University of Iowa.

Departments

Desktop

Keeping Transients and Static in Check

Steven Dettwiler
Product Technical Marketing Engineer
Intel Architecture Group
Intel Corporation

Overview

As processors increase in speed and the available voltage to the processor continues to decrease, transient and static limits are increasingly important to motherboard designs. Intel power delivery designs for desktop motherboards are critical for keeping transients and static in check.

Limiting transients and static in the motherboard helps developers protect the sensitive transistors in the CPU. These limits also help prevent electrical anomalies, which can cause end-user frustration and dissatisfaction. For example, low voltages can cause speed path problems, which can make the CPU run slowly. High voltages can permanently damage the CPU and can also lead to shorter processor lifetimes.

Intel® desktop board VRM (voltage regulation module) solutions comply with strict board layout topologies and include robust components. These boards are designed to meet and very often exceed stringent processor specifications for transients and static. Designs that do not comply with the Intel® VRM specification may not be able to support faster processors or be as robust as those that meet or exceed the VRM requirements.

Transients and Static

Transients are an inherent side effect of a regulator under heavy switching load conditions. Transient voltages are changes in the output of any regulator due to any changes in the load current. The static level is the constant average output voltage after a peak transient event. Basically, transients are voltages that exceed static limits. If not properly limited, transients and static can significantly and detrimentally impact the voltage supplied to the processor.

Design Techniques

Intel uses several design techniques to keep transients in check as required by the Intel VRM spec:

- Adaptive voltage positioning
- Field-Effect Transistors (FET)
- Input/output inductors
- Gate resistors
- Output capacitors

Intel recommends that designers using Intel® processors make sure their motherboards also comply with the Intel VRM spec. (All Intel® motherboard designs comply with this spec.) Motherboards that do not comply with Intel specs may send lower or higher voltage to the processor, causing speed path problems, damage, and shorter lifetimes for the processors.

Voltage Positioning

Adaptive voltage positioning is a technique that reduces the amount of bulk capacitance that the processor requires. The amount of bulk capacitance is reduced, while the output voltage is maintained within its regulation limits during load transients.

This technique positions the output voltage at a higher level under a light load and decreases it under a heavy load. The effect is to pre-charge the capacitors in anticipation of a load transient. Output voltage therefore varies with load current, and thus increases the effective impedance of the converter.

Intel motherboards also include a sense resistor as part of the adaptive circuitry. The sense resistor helps set current limits and takes adaptive design concepts a step further using load lines. The load lines lower the total power dissipation of the processor. This improves the overall manufacturing yield for Intel® Pentium® processors.

FETs

Switching FETs (Field-Effect Transistors) used in Intel designs are based on the motherboard's requirements for three important factors:

- Power dissipation
- Low resistance-drain-source
- Gate charge

FETs are necessary because higher speed processors demand additional current. FETs are also useful where it is critical to have high-frequency processors and efficient power use.

Inductors

Input/output inductors are inexpensive components that can help improve transients. In Intel motherboards, the inductors are usually thermal-mount, single winding, with a toroid core. They help set the output voltage ripple. To some degree, they also limit regulator response.

Input inductors limit the current draw from the VCC board plane. Input inductors also help attenuate noise and EMI (electromagnetic interference), which are generated by the synchronous buck regulator. This is a key design element for helping to keep transients within specifications.

Gate Resistors

Gate resistors are in place to prevent crossover, overshoot, and undershoot during an ON event from the pulse-width-modulation controller. Without gate resistors, momentary shorts can occur across the input bus. These intermittent shorts can result in lower efficiencies. Such shorts can also destroy the switching devices.

Output Capacitors

Output capacitors are low-pass filters that help maintain regulation during load transients and control the output voltage ripple that is induced by the buck regulator switching converter.

Oscons are organic semiconductors, and are the capacitor of choice for controlling transients and static. They are used for both output bulk and filtering. When compared to tantalum or aluminum capacitors, oscons have higher ripple current ratings, low ESR (equivalent series resistance) and low ESL (equivalent series inductance), a smaller package size, and a longer lifetime. Aluminum capacitors also cannot provide the low parasitics required for high-speed processors like the Pentium® 4 processor.

Layout Topologies

Figure 1 shows a typical board layout of the buck regulator used in Intel motherboard designs. Proper layout and component placement are critical to reduce thermal dissipation of high-current FETs, inductors, and sense resistors, and to reduce parasitics.

Figure 1 shows careful consideration to trace flow to help reduce thermal output. This layout avoids neck downs and sharp angles on power planes. The topology also uses MAX copper around all FETs. In some areas, heavy trace widths (note the FET drive gate) help ensure proper thermal management. There are also heavy trace widths (ground path arrow) from ground planes to the nearest capacitor. These trace widths provide proper return paths for parasitics and help in reducing EMI.

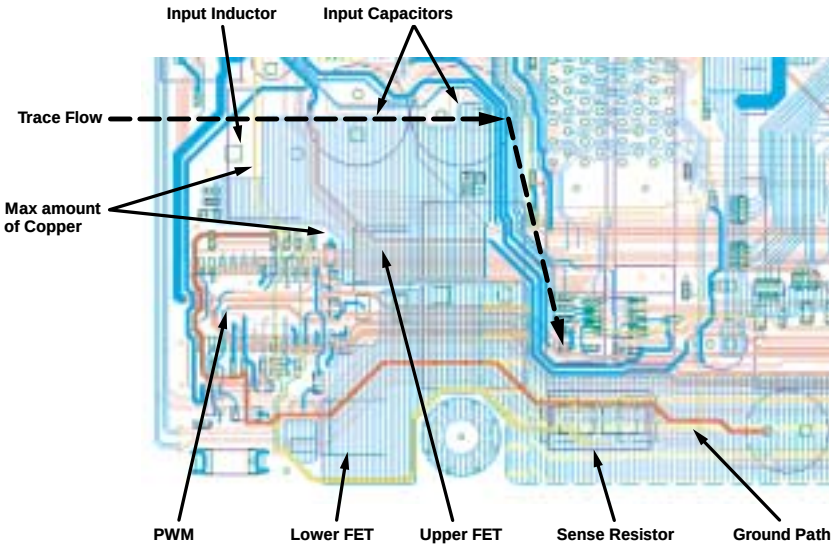

Figure 1. Buck Regulator Layout

Summary

Intel uses a variety of techniques and components to reduce transients and static in today's motherboard designs, including adaptive voltage positioning, FETs, I/O inductors, gate resistors, and output capacitors. These power delivery designs help Intel motherboards protect CPUs and extend processor lifetimes. By using robust components and complying with strict board layout topologies, Intel offers motherboards that meet and often exceed stringent specifications for limiting transients and static in today's desktop systems.

To limit potential damage to the CPU and ensure longer processor lifetimes, make sure that your motherboard complies with the Intel VRM specification.

More Info

For a list of Intel-specified motherboards, check the support pages of the Intel Web site. Articles and white papers on voltage regulation can be found in the desktop hardware section of the Intel Developer site. The Intel VRM 8.4 DC-DC Converter Design Guidelines define a range of DC-to-DC converters to meet the power requirements of computer systems using Intel® microprocessors. The VRM 8.4 guidelines can be found online at the Intel Developer Web site.

Author Bio

Steven Dettwiler joined Intel more than 10 years ago. He began in the motherboard design group, working on projects such as 386Hi. Later, he joined the OEM Platform Solutions Division, where he has worked on TC430HX, AL440LX, AR440BX, and VC820 products, as well as today's DB850GBC. Steven helped in the successful launch of the 4XAGP and RDRAM technologies. He received his B.S. in engineering from the Oregon Institute of Technology.

New Intel® Desktop Boards Provide More Versatility, Less Cost

Dan Daley
Senior Product Marketing Engineer
IAMG Desktop Boards/Systems Marketing
Intel Corporation

Overview

The new, full ATX-sized Intel® desktop board D815EEA2 and the micro-ATX Intel® desktop board D815EFV are the second generation of boards based on the proven Intel® 815E chipset and the new Intel® 815EP chipset. Supporting both the Intel® Pentium® III and Intel® Celeron™ processors, these new desktop boards offer great performance with more features and options than ever before–and do it for less cost.

Both boards provide a stable foundation for system integrators who want to support upcoming technology and higher frequency processors. The boards' high quality, low cost, and versatility allow developers to easily meet more price points in more market segments. Their flexibility, reliability, easy integration, and leading technology make them the best choice for new mainstream and value-priced systems.

Two sizes to fit all

A glance reveals their form factors (Figure 1). The Intel desktop board D815EEA2 is 11.5 by 8.2 inches (standard ATX size), while the Intel desktop board D815EFV, at just 9.6 by 8.2 inches, easily fits micro-ATX installations.

Figure 1. The Intel® D815EEA2 (left) is a full sized, full-featured ATX desktop board. The smaller Intel® D815EFV (right) desktop board is just as capable and designed for micro-ATX installations.

It takes a closer look to see both boards share a full range of capabilities (Figure 2). Each has support for all Intel Pentium III and Intel Celeron processors, four USB ports, SoundMAX* 2.0* 3D audio, and flexible options for graphics and network connections. The primary difference is that the full-sized board has more PCI slots for further expansion–five versus three on the micro-ATX board.

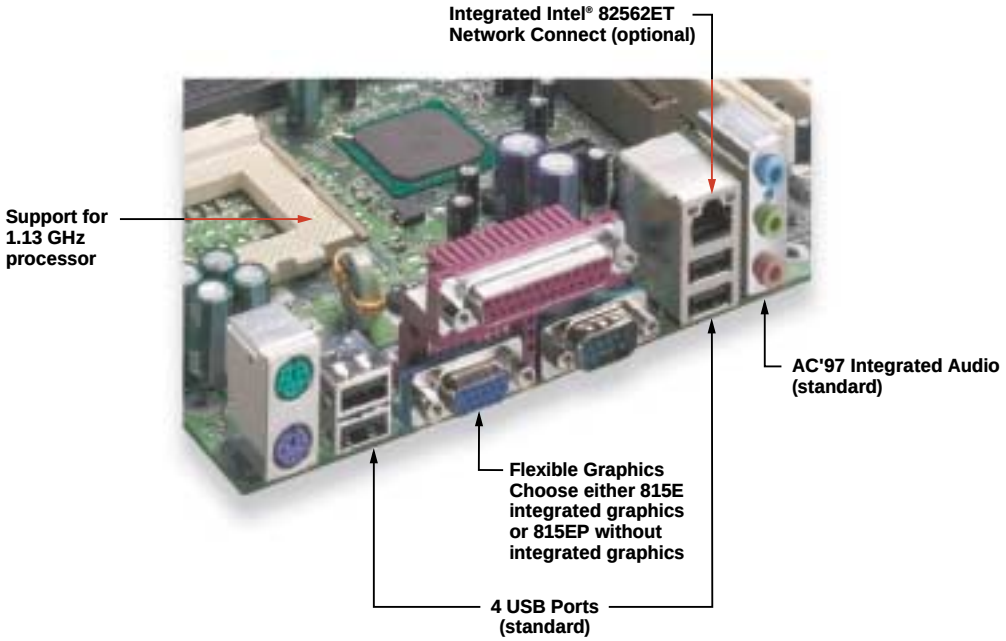

Figure 2. Both form factors (micro and full ATX) share a robust set of capabilities

Integrators can take advantage of the common board design between form factors and leverage software and hardware qualifications. Each board offers options to allow the system integrator to balance graphics, performance, and price points for business and consumer applications. The models without integrated graphics rely on the Intel 815EP chipset, while all other models use the standard Intel 815E chipset with full integrated Intel® 3D Graphics. All versions permit the addition of discrete graphics via a Universal AGP 4X controller.

Rock Solid Reliability

Intel-designed desktop boards can give systems the ability to work right out of the box. Integrators often need to quickly create customized systems. The Intel Compatibility Evaluation Labs have conducted over 500 tests with more than 350 commercially available hardware, network, and software products. These tests lessen the burden on integrators and ensure rock solid compatibility.

This reliability is borne out by the Intel® Active Monitor application included with every Intel desktop board D815EEA2 and D815EFV. This new thermal management utility monitors temperatures, power supply voltages, and fan speeds. It offers constant monitoring and customizable alerts at the first sign of a hardware fault.

Snap Integration

The integration from solid hardware to the end user's fingertips can be long and expensive. "Snap Integration" with Intel® value-added software, another feature of the Intel desktop boards D815EEA2 and D815EFV, nicely closes that gap.

The Intel® Express Installer, which comes on CD with each motherboard, starts off by getting system software, drivers, and support onto the hard disks with 10 simple clicks versus the 35 to 48 steps of less intelligent setups. It employs a quick checklist interface and a one-button install.

- Intel® Rapid BIOS Boot can bring the end user up to Windows ME* in less than 20 seconds. This speed has been achieved by reducing the power-on self-test (POST) time by as much as 60 percent, while maintaining the complete feature set.
- The Intel® Express BIOS Update is a Windows*-based BIOS update utility with an Easy Install Wizard interface to keep the system up-to-date even with today's rapid technology changes.
- Intel desktop boards D815EEA2 and D815EFV come complete with a solid suite of protection comprising Norton Internet Security* with Norton AntiVirus*, Norton Personal Firewall*, Norton Privacy Control*, and Norton Parental Control.

Leading Technology

The Intel desktop boards D815EEA2 and D815EFV support leading-edge technologies to ensure productivity and long-term reliability. They have three dual in-line memory modules (DIMM) sockets that allow up to 512 MB of memory. Each allows Ultra ATA/100 disk support for faster disk access. Both boards support Intel Pentium III processors with 133/100-MHz system buses and Intel Celeron processors with 100/66-MHz system buses.

Numerous options allow both boards to be integrated to fit different price points and user requirements. Integrators can call on an Intel® PRO/100 Network connection for integrated local area network (LAN) or rely on the boards' Communications and Networking Riser (CNR) support to create other low-cost networking and communications solutions. An optional MicroCross* Molex* DVI Link* riser card is available to take advantage of the Digital Video Output support, allowing developers to directly configure systems for flat-panel and high-resolution digital displays. The boards' Universal 4X AGP support keeps a full range of options open for changing graphics requirements.

Summary

The Intel desktop boards D815EEA2 and D815EFV are a second-generation follow up to the successful Intel 815E chipset. They feature the rock solid reliability, snap integration, and leading technology–exactly what system integrators need to hit today's price points and tomorrow's new applications. The boards' form factors, options, and versatility make them ideal platforms for both the mainstream and value markets.

More Info

To find out more about Intel desktop boards and software, visit the Intel Desktop Board page of the Intel Developer Web site.

For more information on the Intel® 815 chipset, visit the Intel 815 Chipset page of the Intel Developer Web site.

For more information on the MicroCross Molex DVI Link riser card, visit the Molex Web site.

Author Bio

Dan Daley has been with Intel for one year and supported the launch of the desktop board D815EEA. He received divisional recognition for his work on the D815EFV cost reduction team. He has both a B.S. degree in mechanical engineering and a M.S. in mechanical engineering from Brigham Young University. Dan also earned an M.B.A. degree from the Marriott School of Management at Brigham Young University.

Dual Load Switching

Steven Dettwiler
Product Technical Marketing Engineer
Intel Architecture Group
Intel Corporation

Overview

Today's desktop systems have stringent ACPI (Advanced Configuration and Power Interface) requirements. These requirements help desktop systems operate properly and robustly in low-power S3 and S5 (suspend to RAM) sleep states.

For ACPI-compliant desktop systems, the on-board voltage regulator (VR) design for motherboards must take into account all aspects of proper power distribution. One important aspect of power distribution is adequate switching of power rails between main and standby voltages. Proper VR switching implementation will help differentiate a robust ACPI-supported motherboard in today's market.

Designing to Sleep States

Intel® desktop boards support three independent sleep states: S0 (normal operation), S3 (suspend to RAM), and S5 (shutdown).

In S0, the motherboard (or other board) is up and running normal operation. In S0, all voltage regulators are in nominal states from the ATX Power supply box discussed in the Dual Load Switching section. In S0, all current loads are being drawn from the main rails (5V, 12V, 3.3V, 5 VSB).

The S3 sleep state is also known as standby, or Suspend to RAM (STR). S3 typically consumes about 10 percent of the total power required for S0. In S3, major components are off, including corresponding voltage regulators. Memory and Wake On rails are powered by standby voltages (5 VSB, 3.3 VSB).

In S5, the desktop system is fully asleep and main power rails are off. Only WOL (Wake on LAN) components receive power. The motherboard must support booting from a WOL event.

Inadequate Power

Most motherboard manufacturers connect sleep-state power rails directly to the standby power rail. Although this method is inexpensive, it can also result in inadequate power during resume (from sleep state) and thus can cause system failures during S3. For example, a fully loaded desktop system could include USB or a communications network riser (CNR). For a fully loaded system, the amount of current required to wake the system from S3 would easily exceed the standby currents available on a typical ATX power supply.

Table 1 shows which voltage rails are supported by each ACPI sleep state. Voltage rails that support S3 and S5 must switch between main and standby power because of the typical limitation of 5 VSB, which is the only available standby voltage generated by the ATX power supply. This voltage limitation creates the problem of satisfying all standby current demands necessary for waking all hardware components from S3 and S5.

Voltage Rail	S0	S3	S5
Vccp	Y	N	N
Vtt	Y	N	N
GM CH Core	Y	N	N
AGP	Y	N	N
ICH Core/Hublink	Y	N	N
Memory Core	Y	Y	N
USB/PS2	Y	Y	N
CNR3.3	Y	Y	Y
PCIVaux	Y	Y	Y
1.8V ICH LAN	Y	Y	Y

Table 1. Voltage Breakdowns

Dual Load Switching

Intel has resolved the power problem in today's designs by creating a dual load switch (DLS). The DLS solution switches power rails between main and standby voltages using N channel and P channel FETs. By switching power rails, Intel's design avoids the voltage limitation of today's typical ATX power supplies.

A typical Intel® motherboard uses five DLSes to ensure a complete solution for all peripherals. Figure 1 shows the logic control signals used in the DLS, the N and P channel nets, and the dual switchable output from 3.3V main power to 3.3 VSB. As Figure 1 shows, the N channel is driven with 12V and the P channel is turned off with 5 VSB. This enables the logic to control which sleep states are supported.

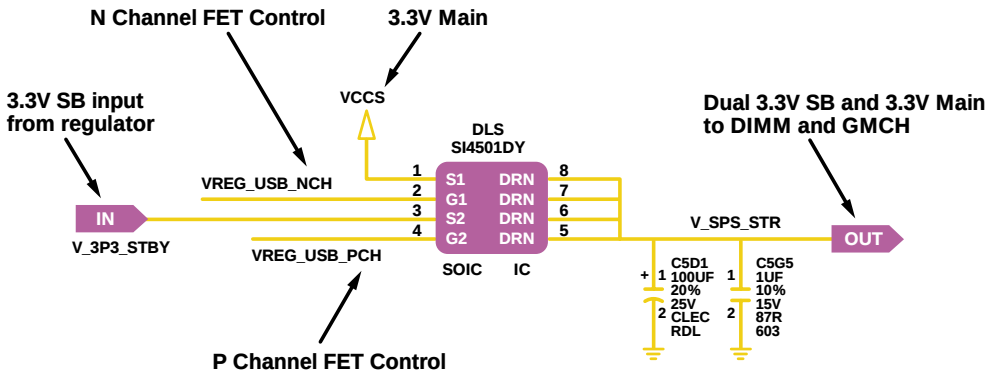

Figure 1. Dual Load Switch (DLS)

Intel's solution ensures that each sleep state is completely supported for each power rail. Both the S3 and S5 sleep states require two dual-load switches for each supported voltage rail. In the Intel solution, in the S3 sleep state, one 3.3V DLS is required for the memory core and one 5V DLS is needed for USB/PS2 devices. In the S5 sleep state, one 3.3V DLS is needed for WOL/PCI Vaux devices, and one 5V DLS is required for communications network riser (CNR) devices.

Summary

One of the ways Intel desktop board power-delivery solutions are differentiated from other OEM boards is through dual load switching. Already in use, DLS is an innovative, cost-effective, and robust ACPI voltage solution to the typical problems of inadequate power during system wake from various sleep states.

All Intel desktop motherboards meet or exceed the requirements of the Intel® voltage-regulation module (VRM) 8.4 DC-DC Converter Design Guidelines. Evaluating your motherboards against the Intel® VRM 8.4 specification will help you validate your OEM boards for proper ACPI 5 and 3.3 standby voltage, which provide adequate power-delivery solutions for your desktop PCs. Motherboards that do not adhere to the VRM specification may not be able to avoid the voltage limitations of typical ATX power supplies that are in use today.

More Info

For information about ACPI and sleep states, refer to the Intel® ACPI Web site. The design guide for instantly available power-managed desktop PCs is provided online at the Intel IAPC site. The VRM 8.4 guidelines can also be found online at the Intel Developer Update Web site.

Author Bio

Steven Dettwiler joined Intel more than 10 years ago. He began in the motherboard design group on projects such as 386Hi. Later, he joined the OEM Platform Solutions Division, where he has worked on TC430HX, AL440LX, AR440BX, and VC820 products, as well as today's DB850GBC. Steven helped in the successful launch of the 4XAGP and RDRAM technologies. He received his B.S. in engineering from the Oregon Institute of Technology.

Initiatives and Technologies

Using the CNR with Multiple Chipset Suppliers

Michael J. Abel
Technical Marketing Engineer
Platform Application Engineering (PAE)
Intel Corporation

Overview

The Communication and Networking Riser (CNR) Specification defines a flexible, cost-effective solution for implementing multiple technologies onto a single riser card. The specification combines the existing interfaces of the Audio Modem Riser (AMR) with additional interfaces for local area network (LAN) and Plug-n-Play (PnP). This supports new and existing technologies such as audio, modem, and LAN/Home Phone Networking Alliance (HPNA), in addition to wireless and broadband technologies such as Bluetooth*, 802.11, cable, and DSL. With the advantage of sharing a PCI slot and PnP capabilities through the System Management Bus (SMBus), the CNR Specification avoids many of the drawbacks of AMR.

Below is a list of some of the benefits the CNR Specification provides:

- Offers flexibility for implementing multiple technologies for specific platform needs.
- Provides a cost-effective solution when using chipsets with integrated controllers (AC'97, LAN controller, USB controller.)
- Reduces motherboard SKUs, SKU management, and inventory control by offering several technologies onto a single riser card.
- Mechanically shares a slot with the PCI connector.
- Reduces noise and interference of sensitive components due to separation from the motherboard of the digital and analog circuits.
- Offers up to 6-channel audio using cost-effective soft audio solution upgrade.

Multiple Chipset Support

The CNR Specification allows system designers to use core logic chipsets not only from Intel, but from several suppliers as well. Since the CNR Specification supports either the MII or the LCI for Ethernet (LAN) or Home Phoneline (HPNA) networking, CNR can be used with different suppliers of core logic chipsets and designed on multiple platforms. For use with non-Intel® chipsets, the MII would be the supported LAN interface and can be discrete or integrated into the chipset as shown below in Figure 1.

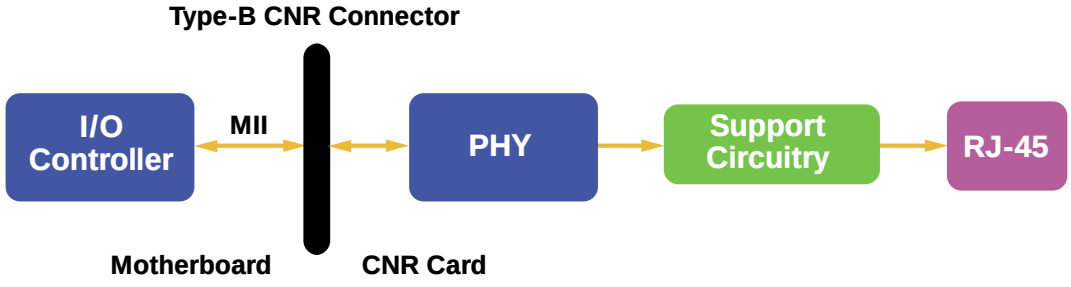

Figure 1. LAN on CNR Using the Media Independent Interface

Figure 1 shows a general implementation of a CNR card for LAN using the MII. The essential components consist of the Type-B connector pinout, used as described in the CNR Specification, and an MII-compliant PHY along with support circuitry consisting of magnetics and an RJ-45 connector. For more detailed information on designing the MII for LAN on CNR, refer to the white paper titled, Recommendations for Implementing the MII on CNR. Multiple chipset suppliers support integrated 10/100 BaseT Ethernet MII-compliant MACs.

Industry Support

Industry support of CNR has been growing tremendously. A list of motherboards using various non-Intel® chipsets with CNR appears in Table 1.

Motherboard Supplier	Product	Chipset
FIC	AD11	AMD761/Via 686
FIC	AZ11E	KT133/Via 686B
FIC	AM33	KM133/Via 686B
Lucky-Star	6AAP5	Ali
Shuttle	AV31	VT8233
Shuttle	AK11	KT133
MSI	MS6340	VIA KT133/686A
MSI	MS6309	694/686B
MSI	MS6321	694XDP/686A
MSI	MS6347	KT133/686A
MSI	K7T Pro2	KT133/686B
MSI	K7T Pro2-A	KT133/686B
ECS	D6VAA	694X/686B
Jetway	849B	Ali

Table 1. Motherboard suppliers using non-Intel chipsets supporting CNR

Summary

The CNR Specification provides a flexible, cost-effective solution for implementing communication and networking technologies with core logic chipsets. Support for CNR continues to grow through industry adoption.

By providing a flexible and cost-effective way to implement audio, modem, LAN/HPNA, wireless, and broadband technologies, CNR strives to be the prominent riser solution. CNR lets integrators take advantage of chipset integration using the AC'97 Link, USB, or LAN interfaces. Since the CNR Specification defines two pinouts, Type-A to be used with the LAN Connect Interface (LCI), or Type-B, which supports the MII implementation, CNR provides the flexibility needed to support multiple chipset suppliers for LAN.

More Info

The Intel® CNR Web site provides press releases, white papers, and marketing material, as well as links to card manufacturers and motherboard suppliers and other information. These pages are of particular interest:

- Communication and Networking Riser (CNR) Specification version 1.1
- List of current motherboard suppliers and card manufacturers
- *Recommendation for Implementing the MII on CNR*, a white paper
- *A Comparison of the New PC Riser Specifications*, a white paper

Author Bio

Michael J. Abel is a technical marketing engineer working in the Intel Architecture Marketing Group. He presented *Implementing Broadband and Wireless on CNR* at the Intel® Developer Forum Conference, Spring 2001 and wrote the Intel white paper titled Recommendations for Implementing the MII on CNR. Michael earned his B.S. in electrical engineering from California Polytechnic State University in San Luis Obispo. He is an active member of the IEEE and is involved in wireless and broadband technologies.

Networking & Communications

Achieving High Availability on Intel® Architecture

Overview

An excerpt from the white paper "Providing Open Architecture High Availability Solutions"

Advanced, complex hardware and software systems have a significant presence in our everyday lives in the tasks and services performed by our telephones, banking institutions, computers, and the Internet. Nearly every industry relies on the availability of large, complex hardware and software systems to perform important roles that range from enhancing productivity in the work-place to providing the technology that advances our search for innovation. As our dependency on complex hardware and software systems increases, the risk and liability that naturally comes with a potential for failure also increases. While some system failures might result in an inconvenience, other system failures could result in loss of revenue.

Definition of High Availability

High Availability, or HA, is the term used to describe computer systems which exhibit almost no downtime. This is typically quantified in terms of the number of "9s," ranging from 3 nines to 6 nines, as shown in Figure 1. In order to provide this type of availability, a system must be designed in a reliable manner, have processes in place to ensure rapid recovery from faults, and must be used only within its design parameters. Additionally, the system must be well managed and secure from unauthorized use and activities.

Number of 9s	Downtime per Year	Typical Application
3 Nines (99.9%)	~ 9 Hours	Typical Desktop or Server
4 Nines (99.99%)	~ 1 Hour	Enterprise Server
5 Nines (99.999%)	~ 5 Minutes	Carrier Class Server
6 Nines (99.9999%)	~ 31 Seconds	Carrier Switch Equipment

Figure 1. High Availability in Terms of the Number of "9s"

Need for an Open Specification

History and the evolution of diverse applications have added complexity and variance to today's HA solutions. Often the development has been ad hoc and narrowly focused, and has led to proprietary solutions that are fragile and expensive to acquire and support. Systems are far too large and complex to expect that a single vendor can provide all of the functionality required. Due to the cost and complexity of developing proprietary systems, telecom and Internet infrastructure equipment manufacturers are quickly realizing the need to develop open architecture building blocks.

Open architecture allows manufacturers of those systems to reduce their time-to-market while maintaining the highest levels of reliability. By providing an open specification framework, interoperable building blocks will perform in a uniform and controlled manner, and the ability to scale or grade the performance and reliability to meet the needs of the system designer will be gained.

In addition, formally managing the participation of these otherwise disparate technologies from multiple vendors will improve reliability. This gives a higher level of confidence in the validation tests that qualify these services before their deployment. Further, an open architecture specification will encourage the development of commercial-off-the-shelf components, foster a competitive environment which will improve product features, cost and performance, and improve the supply of staff with relevant HA technology experience.

Applications That Will Benefit

While the primary application area for an open HA specification is broadly aimed at the systems that provide telecommunications, data communications, and Internet infrastructure, many other application areas share similar types of architectures and requirements, and will also benefit. Some examples of telecommunications and Internet infrastructure applications include:

- SoftSwitches and VoIP gateways
- Telecom network service control points
- Telecom network switch transfer points
- Wireless base station controllers, radio network controllers
- Remote access controllers
- Web hosting, caching, and e-mail servers
- Broadband distribution nodes
- Voice portals and unified messaging systems

These systems are components in the larger network infrastructure. They provide critical services in the data path, control nodes, and network management elements, as well as in the billing and the application service areas.

System Building Blocks

In order to create an open architecture system, interoperable building blocks must be available. These blocks can then be combined as needed to create a system. Major functional blocks are separated into hardware, operating system, management middleware, and applications.

Hardware capabilities of fault managed systems can be generally categorized into three sets:

- Redundancy to allow continued processing after failure
- Highly reliable (often redundant) communication among components
- Platform management, including fault detection, diagnosis, isolation, recovery, and repair

The OS capabilities include functions to isolate, prevent the propagation of, or mask the impact of hardware and software faults. These functions help prevent errant applications or faulted hardware from bringing the entire system down. Another area of capabilities includes dynamic reconfiguration and enhanced device drivers to allow graceful replacement of failed hardware. Additionally, the OS provides services for autonomous fault management, reporting faults externally, and interfacing with HA capabilities in middleware and application layers.

Management middleware is the software component that oversees all of the configuration and fault management services in the system. The availability management component operates automatically in real time and without human intervention. A state-aware system model represents, models, and monitors the status, topology and dependencies of components. System information is collected and assessed, and system anomalies or faults are detected and diagnosed. Faults are acted upon by dynamically reconfiguring the status, configuration, and dependencies of the components to rapidly recover and maintain service. A final capability of management middleware is that it checkpoints (periodically transfers) data between a component and its redundant units in order to maintain operation during system reconfiguration.

Applications used in HA systems can either depend on the HA system to simply restart them if a failure occurs, or they can be “HA-aware.” There are many ways for HA-aware applications to participate, control, or operate within a highly available system. The system should provide a management interface through which an application can monitor operations and send status, heartbeat, and checkpoint information. An application may also need to receive this type of information from other applications. Finally, an application may need to initiate a fail-over or other recovery action and be able to be unloaded, loaded, and restarted while the system is operational.

The HA Forum

The High Availability Forum was created to increase the number and capability of open architecture high availability systems by standardizing the interfaces and capabilities of building blocks for HA systems. The forum intends to promote solid development models and best-known methods for providing high availability systems.

The forum will also consider the environment and full life cycle of components in the target infrastructure. These include 1) compatibility with existing and legacy systems and infrastructure, 2) design, porting, and troubleshooting applications, 3) installation into the infrastructure, 4) manageability and security, 5) performance, cost, and scalability, 6) component failure, isolation, and recovery, 7) testing and repair, and 8) hardware and software changes and upgrades.

The High Availability Forum members are Dialogic, an Intel Company; GoAhead Software; Hewlett-Packard Corporation; Intel Corporation; LynuxWorks, Inc.; MontaVista Software; Motorola Computer Group; RadiSys Corporation; and Ziatech Corporation.

More Info

For more information on High Availability and the High Availability Forum, including the full white paper titled "Providing Open Architecture High Availability Solutions," visit the HA Forum Web site.

Servers

Unisys Takes Open Systems into the Data Center

David L. Houseman
Vice President, Advanced Technology
Systems and Technology Group
Unisys Corporation

Overview

The Cellular MultiProcessing (CMP) architecture establishes an Intel® processor-based high-performance, high availability system capable of supporting large-scale e-Commerce and mission-critical applications. A CMP platform can be configured with up to 32 Intel processors and operated as a single large system or separated into up to eight partitions, each running its own Intel processor and operating system, enabling flexibility and almost linear scalability. The entire unit is easily managed.

CMP technology was designed to enable open systems based on Intel processors with the mission-critical capabilities of a mainframe. Beginning with the particular requirements of database servers, data warehousing, large application servers and ASPs, Unisys Corporation designed a highly available, scalable, and manageable system that supports mission-critical enterprise data center applications.

The performance potential for applications in a 32-processor environment is only beginning to be tapped by software developers. Developers for Intel-based systems should be aware that the CMP platforms are not only shipping today in volume for IA-32 Intel® Architecture processors, but that systems for Intel® Itanium™ processors are expected to ship by summer 2001.

Architecture Based on 32 Intel® Processors

The CMP architecture supports 32 Intel processors in the cabinet, which can be a mixture of IA-32 processors and Itanium processors. These processors are divided into eight hot-pluggable sub-PODs of four processors each. Each sub-POD four-processor grouping has a third-level cache. A crossbar intra-connect allows all of the sub-PODs to communicate with all of the memory, which is composed of up to four Memory Storage Units (MSU) which allow for parallel control and data access. Very high-speed, high bandwidth point-to-point connections are used rather than buses.

A key to a great deal of the parallelism is the third-level cache, with up to 32 MB of external static RAM associated with each one. Given that there can be up to eight of those, a great deal of external cache is associated with the processors. It brings up to 64 gigabytes of main memory much closer to the processors with a 0-wait state access time, increasing both performance and scalability.

In most systems that are bus-based, where the processor must modify or access a cache line, all the caches on the bus must be examined or “snooped,” slowing down performance. With CMP, only one other processor on a bus is snooped, and the rest are handled by directory searches within the memory. This enables up to 133 million cache bindings per second.

The CMP architecture can be partitioned into multiple running operating systems concurrently, including Microsoft Windows* 2000, Microsoft NT* 4.0, and Unixware. Any operating system that will run in an Intel space to MP-compliant specifications will run.

System Scalability and Architectural Scalability

System scalability is the ability to grow a client's applications as business needs change. CMP excels at this sort of scalability. For instance, as database accesses increase and requires more processors, the CMP platform provides the system scalability needed to support that growth.

Architectural scalability is the ability to support multiple types of work loads, such as online transaction processing, data warehousing, data marts, Web traffic, and other types of high-performance computing. The CMP platform was designed to ensure that all the pipes, buses, and other components would support the target design. Modeling was run for a number of different work loads, ensuring clients would have the ability to buy one system flexible enough to run whatever work loads they needed, concurrently, in the same platform.

The CMP platform can be partitioned into four 8-way servers, eight 4-way servers, two 16-way, an 8-way and two 4-ways, or whatever combination in groups of four processors that is needed. It supports multiple processors, including Intel Itanium processors, even concurrently, as well as multiple speeds. This provides great scaling flexibility as applications are developed and as new Intel processors come online.

IMS Provides Manageability

The CMP architecture has an integrated management system (IMS) with stand-alone hardware that can be made available in a redundant configuration. It is used to boot the system, load the hardware, create partitions, and otherwise manage the system. The system is managed from IMS from an overall operations point of view. It comes with a suite of custom applications based on Microsoft Windows that allows for improved manageability throughout the enterprise.

The IMS also has a Web client, enabling remote access capabilities to operations and support staff. This provides a mainframe class of access and management that is new for Intel processor-based products. The Web client can run anywhere and perform any function; for instance, someone at home can do anything from power up from a power-off state to booting an operating system, loading applications, running applications, and operating the multiple operating system environments.

Security

Although security is generally handled by operating systems, the CMP platform has security features built in to the hardware. If a CMP system is soft-partitioned into different running environments, range checking is done at the hardware level on all references to memory. This ensures that an operating system can access only memory to which it has been assigned rights and cannot interfere with another operating system. For instance, if a hacker were to get past one operating system’s security, the hardware would prevent him from interfering with any other operating system in that environment.

Intel-Based Means Value

Compared to other high-end vendors, CMP systems have a relatively low price because they are Intel processor-based and can take advantage of Intel’s volumes. A fully configured CMP system has a street price of about one third of other vendors’ comparable RISC-based systems.

Summary

The CMP architecture enables up to 32 Intel processors to run in the same box, with a single operating environment or multiple environments, and communicate through a very fast shared pathway instead of via TCP/IP or other means, thereby providing a very powerful and unique capability.

More Info

To find more information about CMP architecture for Intel processor-based servers, visit the Powered by CMP Web site.

To find more information about the Unisys ES7000 series, based upon CMP architecture, visit the Unisys Web site.

Author Bio

David Houseman is the chief technologist for the Systems and Technology Group, providing strategic leadership for Unisys technology directions. He also chairs the corporate Unisys Architecture Committee to ensure consistent architecture and technology directions for the company. Additionally his organization has responsibility as the focal point for standards participation, advanced platform development, key technology partnerships, venturing of emerging technology and communications of the Unisys technology message inside and outside the engineering community. David has held a number of management positions in hardware and software product development. He earned his B.S. in electrical engineering from the University of Virginia and an M.S. in electrical engineering from the University of Florida.

—End of Intel Developer Update Magazine Issue 19—